

BRIEF REPORT

Creating Illusions of Knowledge: Learning Errors That Contradict Prior Knowledge

Lisa K. Fazio
Carnegie Mellon University

Sarah J. Barber
University of Southern California

Suparna Rajaram
Stony Brook University

Peter A. Ornstein
University of North Carolina at Chapel Hill

Elizabeth J. Marsh
Duke University

Most people know that the Pacific is the largest ocean on Earth and that Edison invented the light bulb. Our question is whether this knowledge is stable, or if people will incorporate errors into their knowledge bases, even if they have the correct knowledge stored in memory. To test this, we asked participants general-knowledge questions 2 weeks before they read stories that contained errors (e.g., “Franklin invented the light bulb”). On a later general-knowledge test, participants reproduced story errors despite previously answering the questions correctly. This misinformation effect was found even for questions that were answered correctly on the initial test with the highest level of confidence. Furthermore, prior knowledge offered no protection against errors entering the knowledge base; the misinformation effect was equivalent for previously known and unknown facts. Errors can enter the knowledge base even when learners have the knowledge necessary to catch the errors.

Keywords: fiction, false memory, suggestibility, knowledge

Most people know that the Pacific is the largest ocean on Earth, that Edison invented the light bulb, and that the cheetah is the fastest land animal. Is this knowledge stable, or can it be easily altered? The literature on false memory creation contains many demonstrations of people misremembering details of events (e.g., Loftus, 1975). People can even falsely recall entire events that never occurred (e.g., Loftus & Pickrell, 1995; Wade, Garry, Read, & Lindsay, 2002), with consequences for later behavior (e.g., Geraerts et al., 2008). Although much is known about modifying memory for one’s personal past (e.g., Garry & Gerrie, 2005; Loftus, 2004; Roediger & McDermott, 2000), less is understood

about modifying one’s knowledge base. Our question is whether general knowledge about the world can be easily changed, with people coming to believe that Franklin invented the light bulb or that the Atlantic is the largest ocean.

One’s knowledge base may be more resistant to change than personal memories. These two types of memories depend upon different brain areas (Prince, Tsukiura, & Cabeza, 2007), and can be dissociated behaviorally (Tulving, 1985). In addition, although most episodic memories refer to unique events, most basic facts have been encountered repeatedly, increasing their strength in memory. Facts are also associated with numerous other facts; in addition to knowing that Edison invented the light bulb, people know that he was an American inventor who lived in the 1800s. This supporting information helps to stabilize one’s knowledge about Edison and his invention (Myers, O’Brien, Balota, & Toyofuku, 1984). In fact, research on naive science beliefs suggests that prior knowledge can be extremely resistant to change; even college-level physics courses fail to correct many students’ naive misconceptions. (Carey, 1986; Clement, 1982; Posner, Strike, Hewson, & Gertzog, 1982).

It is unsurprising when people learn factual errors in novel domains. Without prior knowledge, this simply involves encoding new memories according to the principles of episodic memory. More interesting is whether people learn errors despite having the correct knowledge stored in memory. Several studies hint at this

This article was published Online First May 21, 2012.

Lisa K. Fazio, Department of Psychology, Carnegie Mellon University; Sarah J. Barber, Davis School of Gerontology, University of Southern California; Suparna Rajaram, Department of Psychology, Stony Brook University; Peter A. Ornstein, Department of Psychology, University of North Carolina at Chapel Hill; Elizabeth J. Marsh, Department of Psychology and Neuroscience, Duke University.

A collaborative activity award from the James S. McDonnell Foundation supported this work. The opinions expressed are those of the authors and do not represent the views of the foundation.

Correspondence concerning this article should be addressed to Lisa K. Fazio, Department of Psychology, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213. E-mail: lkfazio@cmu.edu

possibility. For example, reading factual inaccuracies in stories (e.g., Marsh, Meade, & Roediger, 2003) or questions (Bottoms, Eslick, & Marsh, 2010) increases the likelihood that people will answer later general-knowledge questions with those errors. However, rather than by directly measuring prior knowledge in an individual, such knowledge has been inferred from norms (Nelson & Narens, 1980) or a postexperiment knowledge check. This study breaks new ground by directly measuring individuals' knowledge 2 weeks before they read stories containing errors.

To establish what individuals knew, we first asked participants a series of general-knowledge questions as part of an online survey purportedly about tip-of-the-tongue states. Then, at least 2 weeks later, the participants completed a seemingly unrelated experiment in the laboratory. During this session, participants read two short fictional stories that contained incorrect information (e.g., "Newton proposed the theory of relativity"). The stories were clearly labeled as fictional, and readers were warned that the stories might contain errors. Soon after reading the stories, participants took a final general-knowledge test on which they were warned against guessing; a subset of these questions referred to the facts mentioned in the stories (e.g., "Who proposed the theory of relativity?"). Of interest was whether suggestibility depended upon individuals' previously demonstrated knowledge. We expected that participants would answer general-knowledge questions with story errors if they were unable to answer the questions correctly on the survey; of primary interest was whether participants would answer with story errors even after demonstrating that they had known the correct answers on the earlier survey.

Method

Participants

Twenty-four undergraduates from the University of North Carolina at Chapel Hill participated for course credit. Two participants were eliminated because they did not answer at least two questions correctly in each set of questions on the online survey, leaving 22 participants in the analyses.

Materials

To measure prior knowledge, we used an online survey containing 64 general-knowledge short-answer questions (from Nelson & Narens, 1980). Sample questions included "What is the largest ocean in the world?" and "What is the last name of the man who invented the telegraph?" The questions were selected to have a range of difficulty. Thirty-two questions measured knowledge of critical facts that were later referenced in the stories, and the rest were filler questions. The 32 critical questions were split into two sets, which were matched for difficulty according to the Nelson and Narens (1980) norms. One set appeared later in the stories as misleading items, and the other set appeared as neutral references. The sets rotated across participants, so items that some participants saw as misleading, others saw as neutral.

Two fictional stories (modified from Marsh, 2004) were used, each of which contained character, dialogue, and plot. The stories were approximately 1,400 words in length. Critically, each story made eight references to incorrect but plausible answers (misleading items) and eight references to critical concepts without sug-

gesting a specific answer (neutral references). For example, a misleading reference referred to "paddling around the largest ocean, the Atlantic," whereas the neutral reference mentioned "paddling around the largest ocean" without specifying the name of the ocean. Across the two stories there were 32 critical facts, which corresponded to the 32 critical questions on the online survey. Each story also included eight correct filler facts (which differed from the fillers on the survey), so that accurate information was also presented. These correct fillers were the same for all participants; they did not overlap with the 32 critical facts and were not considered part of the experimental design. The final general-knowledge test contained 64 short-answer questions: the 32 critical questions from the online survey (also referred to in the stories) along with 32 new fillers.

Procedure

Participants completed an online survey, which was described as investigating tip-of-the-tongue states. They were warned that some of the questions would be difficult and told to respond with "I don't know" instead of guessing. Participants rated their confidence in each answer using a 5-point scale. Later, participants were invited to participate in a seemingly unrelated experiment that occurred at least 2 weeks later in the laboratory ($M = 19$ days, $SD = 4$).

The laboratory session began with a story-reading phase. Participants were presented with a booklet that contained both stories and were told to read the stories carefully, as they would later be asked questions about the stories. Participants were warned that the stories were fictional and that some of the information presented might be incorrect. Following each story, participants answered four comprehension questions. After reading both stories at their own pace, participants solved puzzles for 5 min before taking the final general-knowledge test, with instructions warning against guessing. Finally, participants read the correct version of each fact and were debriefed about the experiment.

Results

Online Survey

Participants correctly answered 38% of the critical questions. Questions that later appeared as neutral facts ($M = .38$) and misleading facts ($M = .39$) were equally difficult ($t < 1$).

Final Test

Participants' responses on the final test were coded as correct, misinformation, other wrong, and "don't know." For the question about the largest ocean, "Pacific" would be scored as correct, "Atlantic" as misinformation, and "Indian" as another wrong answer. Overall, participants answered "don't know" to 42% of questions, indicating that they followed the warning against guessing on the final test.

Of primary interest were the effects of reading errors that contradicted previously demonstrated knowledge. Thus, we first analyzed only items that were answered correctly on the online survey.

The critical result involved ability to answer general-knowledge questions after reading neutral versus misleading references in the stories. As shown in Figure 1A, participants answered fewer questions correctly after reading the misinformation (.82 vs. .71), $t(21) = 2.23$, $SEM = 0.05$, $d = 0.47$, even though they had answered all of these questions correctly 2 weeks earlier. Numerically, a similar pattern occurred when the analysis was restricted to the strongest prior beliefs, which were defined as correct survey answers produced with the highest level of confidence (.95 vs. .87), $t(16) = 1.22$, $SEM = 0.05$, $p = .24$, $d = 0.46$.

Most importantly, participants answered more general-knowledge questions with misinformation after reading it in the stories (.01 vs. .20), $t(21) = 3.52$, $SEM = 0.05$, $d = 1.12$, as shown in Figure 1B. Reading the misinformation was harmful, even when the analysis was restricted to the strongest prior beliefs (.00 vs. .11), $t(16) = 2.21$, $SEM = 0.05$, $d = 0.76$.

We also examined the effects of story reading when participants did not have accurate prior knowledge. The analysis below includes only questions that were answered incorrectly or with “I don’t know” on the online survey (the results were very similar when incorrect and “don’t know” results were considered separately). As shown in Table 1, participants were unlikely to correctly answer final general-knowledge questions if they did not know the answers 2 weeks earlier ($M = .06$), and this was unaffected by reading neutral versus misleading information ($t < 1$). This finding is not surprising, since participants never saw these correct answers in the experiment. In contrast, when participants did not have accurate prior knowledge, there was a large misinformation effect. Reading story errors greatly increased the likelihood of responding with misinformation on the general-knowledge test, as compared with having read neutral references, $t(21) = 5.17$, $SEM = 0.04$, $d = 1.36$.

In the final analysis, we examined whether suggestibility differed for facts for which participants had previously demonstrated knowledge versus ignorance. In order to make this comparison, we needed to adjust for the different base rates, since participants were more likely to respond with misinformation after reading neutral references for unknown facts (.08, as shown in Table 1) than for known facts (.01, as shown in Figure 1B). To do this, we took the

Table 1

Questions on the Final General-Knowledge Test Answered Correctly or With Misinformation Given That Students Failed to Demonstrate the Requisite Knowledge on the Online Survey (and Instead Answered the Relevant Questions Incorrectly or With “Don’t Know”)

Questions answered	Neutral frame	Misleading frame
Proportion correct	.06	.06
Proportion misinformation	.08	.26

proportion of questions answered with misinformation after reading the misleading frame and subtracted the proportion answered with misinformation following the neutral frame. Interestingly, the misinformation effect was no larger when participants did not have accurate prior knowledge ($M = .18$) than when they had been able to answer the survey questions correctly ($M = .19$), $t < 1$. Prior knowledge offered participants no protection from learning the errors presented in the stories.

Discussion

Reading stories containing misinformation led participants to reproduce factual inaccuracies that contradicted their previously demonstrated knowledge. These errors entered the knowledge base after a single exposure, even though participants were warned that the stories were fictional and contained errors, and were further warned not to guess on the final general-knowledge test. Given that they had produced the correct answers 2 weeks earlier, the errors should have been blatant to readers. Furthermore, readers only needed to recognize the errors to avoid them, as opposed to recalling the correct answers. Even so, reading the errors increased their later production, in contrast to the typical episodic memory finding where blatant errors act as warnings and reduce suggestibility (e.g., Loftus, 1979). This misinformation effect was just as big as the one observed when participants did not have any relevant prior knowledge. Having the knowledge necessary to catch or correct the errors did not prevent them from entering the knowledge base.

These results provide further support for the claim that episodic false memories and illusions of knowledge are different from one another. More generally, many manipulations known to reduce memory errors when remembering specific events either have no effect or even increase illusions of knowledge. For example, in eyewitness and other episodic memory paradigms, preencoding warnings, slowed presentation, source monitoring instructions, and blatant errors all reduce suggestibility (Gallo, McDermott, Percer, & Roediger, 2001; Greene, Flynn, & Loftus, 1982; Lindsay & Johnson, 1989; Loftus, 1979; Tousignant, Hall, & Loftus, 1986), whereas these manipulations either have no effect (Marsh & Fazio, 2006; Marsh et al., 2003) or increase learning of errors that contradict general knowledge (Fazio & Marsh, 2008b). In addition, populations such as older adults and children typically make more episodic memory errors, but are less likely to show illusions of knowledge (Fazio & Marsh, 2008a; Marsh, Balota, & Roediger, 2005).

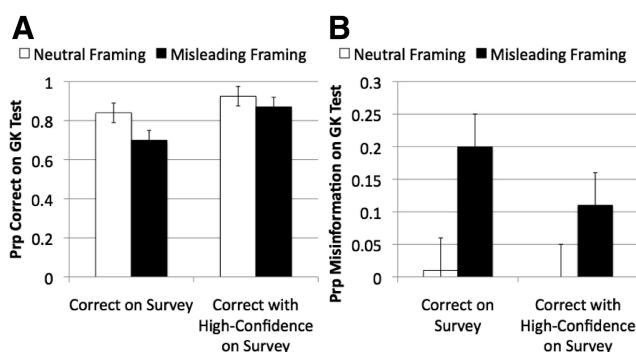


Figure 1. Figure 1A shows the proportion of final general-knowledge questions answered correctly after knowledge was demonstrated on the online survey. Figure 1B shows the proportion of final general-knowledge questions answered with misinformation after knowledge was demonstrated on the online survey. Error bars represent standard errors. Prp = proportion; GK = general knowledge.

Our results support the idea that illusions of knowledge are at least partly driven by *knowledge neglect*: Participants have the requisite knowledge stored in memory, but fail to bring it to bear when it should be used. In line with this idea, when readers are explicitly asked to mark errors in stories, they miss most of them (Marsh & Fazio, 2006). This problem is not limited to the story-reading paradigm used here; students often fail to notice inconsistencies in general-knowledge questions (Bottoms et al., 2010), and readers sometimes fail to benefit from activating related world knowledge prior to reading historical discrepancies (Rapp, 2008). Undetected errors accrue fluency and later come to mind at test, and this retrieval fluency is interpreted as truth (Kelley & Lindsay, 1993; Schwarz, Sanna, Skurnik, & Yoon, 2007).

Note that this account does not involve overwriting of the original memories; we do not believe that the error replaces the correct knowledge in memory, but rather that the two representations coexist in memory. The error was encountered much more recently, however, and thus is more accessible at test. Consistent with this account, manipulations that increase attention to the errors (and thus their later accessibility) increase suggestibility; slowing presentation and highlighting errors both increase the chance that readers will encode them, increasing the likelihood that they will later come to mind fluently (Eslick, Fazio, & Marsh, 2011; Fazio & Marsh, 2008b; Marsh et al., 2003). As this activation fades (e.g., over time), so does the likelihood that participants will produce the errors on later general-knowledge tests (Barber, Rajaram, & Marsh, 2008; Marsh et al., 2003). The fading of activation is not surprising; what is surprising is that a single recent exposure to an error increases its accessibility above that of a strongly held prior response. Anything that further strengthens that activation (e.g., retrieval practice) will make it more likely that the error will continue to be produced after a delay (Barber et al., 2008).

Practically, our results show that educators need to be careful when teaching with fiction or other sources that might contain errors (e.g., interviews of politicians, feature films, television shows, and even some documentaries). Such approaches are common practice, since students generally find nontraditional sources more interesting (e.g., Fehim Kennedy, Şenses, & Ayan, 2011; Fernald, 1987), and yet this approach may lead students to learn errors, even if they already possess the correct knowledge. Educators need to take extra steps and actually point out the exact errors to reduce the harmful effects of exposure to errors (Butler, Zaromb, Lyle, & Roediger, 2009). Similar problems exist in advertising and politics, where consumers and constituents are frequently exposed to false claims that may affect their beliefs, even if they contradict prior knowledge. The more general point is that people's knowledge about the world is malleable and errors can easily enter the knowledge base, regardless of what is already stored in memory.

References

- Barber, S. J., Rajaram, S., & Marsh, E. J. (2008). Fact learning: How information accuracy, delay and repeated testing change retention and retrieval experience. *Memory, 16*, 934–946. doi:10.1080/09658210802360603
- Bottoms, H. C., Eslick, A. N., & Marsh, E. J. (2010). Memory and the Moses illusion: Failures to detect contradictions with stored knowledge yield negative memorial consequences. *Memory, 18*, 670–678. doi:10.1080/09658211.2010.501558
- Butler, A. C., Zaromb, F. M., Lyle, K. B., & Roediger, H. L., III. (2009). Using popular films to enhance classroom learning: The good, the bad, and the interesting. *Psychological Science, 20*, 1161–1168. doi:10.1111/j.1467-9280.2009.02410.x
- Carey, S. (1986). Cognitive science and science education. *American Psychologist, 41*, 1123–1130. doi:10.1037/0003-066X.41.10.1123
- Clement, J. (1982). Students' preconceptions in introductory mechanics. *American Journal of Physics, 50*, 66–71. doi:10.1119/1.12989
- Eslick, A. N., Fazio, L. K., & Marsh, E. J. (2011). Ironic effects of drawing attention to story errors. *Memory, 19*, 184–191. doi:10.1080/09658211.2010.543908
- Fazio, L. K., & Marsh, E. J. (2008a). Older, not younger, children learn more false facts from stories. *Cognition, 106*, 1081–1089. doi:10.1016/j.cognition.2007.04.012
- Fazio, L. K., & Marsh, E. J. (2008b). Slowing presentation speed increases illusions of knowledge. *Psychonomic Bulletin & Review, 15*, 180–185. doi:10.3758/PBR.15.1.180
- Fehim Kennedy, N., Şenses, N., & Ayan, P. (2011). Grasping the social through movies. *Teaching in Higher Education, 16*, 1–14. doi:10.1080/13562517.2010.507305
- Fernald, L. D. (1987). Of windmills and rope dancing: The instructional value of narrative structures. *Teaching of Psychology, 14*, 214–216. doi:10.1207/s15328023top1404_5
- Gallo, D. A., McDermott, K. B., Percer, J. M., & Roediger, H. L., III. (2001). Modality effects in false recall and false recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27*, 339–353. doi:10.1037/0278-7393.27.2.339
- Garry, M., & Gerrie, M. P. (2005). When photographs create false memories. *Current Directions in Psychological Science, 14*, 321–324. doi:10.1111/j.0963-7214.2005.00390.x
- Geraerts, E., Bernstein, D. M., Merckelbach, H., Linders, C., Raymaekers, L., & Loftus, E. F. (2008). Lasting false beliefs and their behavioral consequences. *Psychological Science, 19*, 749–753. doi:10.1111/j.1467-9280.2008.02151.x
- Greene, E., Flynn, M. S., & Loftus, E. F. (1982). Inducing resistance to misleading information. *Journal of Verbal Learning and Verbal Behavior, 21*, 207–219. doi:10.1016/S0022-5371(82)90571-0
- Kelley, C. M., & Lindsay, D. S. (1993). Remembering mistaken for knowing: Ease of retrieval as a basis for confidence in answers to general knowledge questions. *Journal of Memory and Language, 32*, 1–24. doi:10.1006/jmla.1993.1001
- Lindsay, D. S., & Johnson, M. K. (1989). The eyewitness suggestibility effect and memory for source. *Memory & Cognition, 17*, 349–358. doi:10.3758/BF03198473
- Loftus, E. F. (1975). Leading questions and the eyewitness report. *Cognitive Psychology, 7*, 560–572. doi:10.1016/0010-0285(75)90023-7
- Loftus, E. F. (1979). Reactions to blatantly contradictory information. *Memory & Cognition, 7*, 368–374. doi:10.3758/BF03196941
- Loftus, E. F. (2004). Memories of things unseen. *Current Directions in Psychological Science, 13*, 145–147. doi:10.1111/j.0963-7214.2004.00294.x
- Loftus, E. F., & Pickrell, J. E. (1995). The formation of false memories. *Psychiatric Annals, 25*, 720–725.
- Marsh, E. J. (2004). Story stimuli for creating false beliefs about the world. *Behavior Research Methods, Instruments, & Computers, 36*, 650–655. doi:10.3758/BF03206546
- Marsh, E. J., Balota, D. A., & Roediger, H. L., III. (2005). Learning facts from fiction: Effects of healthy aging and early-stage dementia of the Alzheimer type. *Neuropsychology, 19*, 115–129. doi:10.1037/0894-4105.19.1.115
- Marsh, E. J., & Fazio, L. K. (2006). Learning errors from fiction: Diffi-

- culties in reducing reliance on fictional stories. *Memory & Cognition*, 34, 1140–1149. doi:10.3758/BF03193260
- Marsh, E. J., Meade, M. L., & Roediger, H. L., III. (2003). Learning facts from fiction. *Journal of Memory and Language*, 49, 519–536. doi:10.1016/S0749-596X(03)00092-5
- Myers, J. L., O'Brien, E. J., Balota, D. A., & Toyofuku, M. L. (1984). Memory search without interference: The role of integration. *Cognitive Psychology*, 16, 217–242. doi:10.1016/0010-0285(84)90008-2
- Nelson, T. O., & Narens, L. (1980). Norms of 300 general-information questions: Accuracy of recall, latency of recall, and feeling-of-knowledge ratings. *Journal of Verbal Learning and Verbal Behavior*, 19, 338–368. doi:10.1016/S0022-5371(80)90266-2
- Posner, G. J., Strike, K. A., Hewson, P. W., & Gertzog, W. A. (1982). Accommodation of a scientific conception: Toward a theory of conceptual change. *Science Education*, 66, 211–227. doi:10.1002/sci.3730660207
- Prince, S. E., Tsukiura, T., & Cabeza, R. (2007). Distinguishing the neural correlates of episodic memory encoding and semantic memory retrieval. *Psychological Science*, 18, 144–151. doi:10.1111/j.1467-9280.2007.01864.x
- Rapp, D. N. (2008). How do readers handle incorrect information during reading? *Memory & Cognition*, 36, 688–701. doi:10.3758/MC.36.3.688
- Roediger, H. L., III, & McDermott, K. B. (2000). Tricks of memory. *Current Directions in Psychological Science*, 9, 123–127. doi:10.1111/1467-8721.00075
- Schwarz, N., Sanna, L. J., Skurnik, I., & Yoon, C. (2007). Metacognitive experiences and the intricacies of setting people straight: Implications for debiasing and public information campaigns. *Advances in Experimental Social Psychology*, 39, 127–161. doi:10.1016/S0065-2601(06)39003-X
- Tousignant, J. P., Hall, D., & Loftus, E. F. (1986). Discrepancy detection and vulnerability to misleading postevent information. *Memory & Cognition*, 14, 329–338. doi:10.3758/BF03202511
- Tulving, E. (1985). How many memory systems are there? *American Psychologist*, 40, 385–398. doi:10.1037/0003-066X.40.4.385
- Wade, K. A., Garry, M., Read, J. D., & Lindsay, D. S. (2002). A picture is worth a thousand lies: Using false photographs to create false childhood memories. *Psychonomic Bulletin & Review*, 9, 597–603. doi:10.3758/BF03196318

Received October 27, 2011

Revision received April 11, 2012

Accepted April 13, 2012 ■